Parallel Local Algorithms for Core, Truss, and Nucleus Decompositions

Ahmet Erdem Sarıyüce
Sandia National Laboratories
Livermore, CA
asariyu@sandia.gov

C. Seshadhri
University of California
Santa Cruz, CA
sesh@ucsc.edu

Ali Pinar
Sandia National Laboratories
Livermore, CA
apinar@sandia.gov

arXiv:1704.00386v1 [cs.SI] 2 Apr 2017

ABSTRACT

Finding the dense regions of a graph and relations among them is a fundamental task in network analysis. Nucleus decomposition is a principled framework of algorithms that generalizes the k-core and k-truss decompositions. It can leverage the higher-order structures to locate the dense subgraphs with hierarchical relations. Computation of the nucleus decomposition is performed in multiple steps, known as the peeling process, and it requires global information about the graph at any time. This prevents the scalable parallelization of the computation. Also, it is not possible to compute approximate and fast results by the peeling process, because it does not produce the densest regions until the algorithm is complete. In a previous work, Lu et al. proposed to iteratively compute the h-indices of vertex degrees to obtain the core numbers and prove that the convergence is obtained after a finite number of iterations. In this work, we generalize the iterative h-index computation for *any* nucleus decomposition and prove convergence bounds. We present a framework of local algorithms to obtain the exact and approximate nucleus decompositions. Our algorithms are pleasingly parallel and can provide approximations to explore time and quality trade-offs. Our shared-memory implementation verifies the efficiency, scalability, and effectiveness of our algorithms on real-world networks. In particular, using 24 threads, we obtain up to 4.04x and 7.98x speedups for k-truss and $(3,4)$ nucleus decompositions.

Keywords

k-core, k-truss, nucleus decomposition, local algorithms, parallel dense subgraph discovery

1. INTRODUCTION

One of the characteristic features of the real-world graphs is the sparsity at the global level, and the density in the local neighborhoods [12]. Dense subgraphs are strong indicators for functional units or unusual behaviors. They have been helpful in various applications, such as detecting the DNA motifs in biological networks [9], identifying the news stories from microblogging streams in real-time [2], finding the price value motifs in financial networks [8], and locating the spam link farms in web [18, 10, 7], to name a few. Dense regions are also used to improve the efficiency of computation-heavy tasks like distance query computation [17] and materialized per-user view creation [11].

Identifying the significant or anomalous parts of a graph can reveal further insights that are hidden in its complex structure. For this reason, detecting dense structures in various granularities and finding the hierarchical relations among them is a fundamental problem in graph mining. Considering a paper citation network, for instance, investigating the hierarchical relations of dense parts in various granularities can reveal how the new research areas are initiated or which research subjects became popular in time [29]. k-core [30, 21] and k-truss decompositions [27, 5, 33, 34] are two effective ways to find many dense regions in a graph and construct a hierarchy among them. k-core is based on the vertices and their degrees, whereas k-truss relies on the edges and their triangle counts.

Higher-order structures, also known as motifs or graphlets, have been used to locate dense regions that cannot be detected otherwise with edge-centric methods [4, 32]. Finding the frequency and distribution of triangles and other small motifs in real-world networks is a simple yet effective approach used in data analysis [15, 23, 1, 26]. Nucleus decomposition is a framework of decompositions that is able to use higher-order structures to find dense subgraphs with hierarchical relations [28, 29]. It generalizes the k-core and k-truss approaches and finds higher-quality dense subgraphs with more detailed hierarchies. However, existing algorithms in the nucleus decomposition framework require global graph information at each step and can be computationally unaffordable for massive networks. They are also unamenable for parallelization or approximation due to their interdependent iterative nature. We introduce a framework of algorithms for the nucleus decomposition that uses only the local information. Our algorithms are pleasingly parallel and can also be used to achieve faster and approximate solutions to the nucleus decomposition.

1.1 Problem and Challenges

The standard method to compute a k-core decomposition is a sequential algorithm, known as the peeling process. To find a k-core, all the vertices with degree less than k are removed repeatedly until no such vertex remains. This process is repeated after incrementing k until no vertices remain. Batagelj and Zaversnik introduced a bucket-based

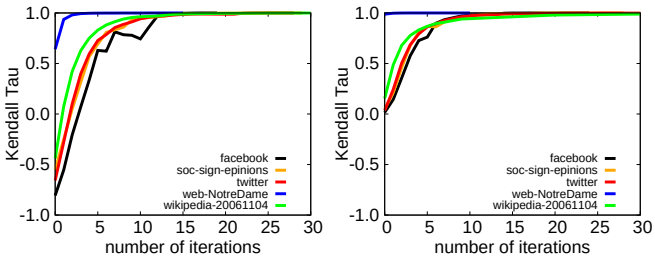

(a) ***k*-truss** (b) **(3, 4)-nucleus**

Figure 1: **Convergence rates with the number of iterations for five graphs in our dataset. Kendall-Tau similarity score compares the obtained decomposition and the exact decomposition; becomes 1.0 when both are the same. We observe that our local algorithms compute almost-exact decompositions in around 10 iterations for *k*-truss, and (3, 4) nucleus decompositions. This enables exploring trade-offs between time and quality, which are necessary for the real-world applications.**

$O(|E|)$ algorithm for this process [3]. It keeps track of the vertex with the minimum degree at each step, thus requires global information about the graph at any time. k-truss decomposition has a similar peeling process with $O(|\triangle|)$ complexity [5]. To find a k-truss, all the edges with less than k triangles are removed recursively and at each step, algorithm keeps track of the edge with the minimum triangle count, which requires information from all around the graph. Nucleus decomposition [28] also facilitates the peeling process on the given higher-order structures. **The computational bottleneck in the peeling process is the need for the global graph information.** This results in inherently sequential processing. Parallelizing the peeling process in a scalable way is not possible since each step depends on the results of the previous step. Parallelizing each step in itself is also infeasible since synchronizations are needed to decrease the degrees of the vertices which are the neighbor to the multiple vertices being processed in that step.

Iterative *h*-index computation: Lu et al. introduced an alternative formulation for k-core decomposition [20]. They proposed to iteratively compute h-indices on the vertex degrees to find the core numbers of vertices. Degrees of the vertices are found at the beginning and each vertex computes the h-index value for the list of its neighbors' degrees. This process is repeated on these values until convergence. At the end, each vertex has the value of its core number. They prove that convergence of degrees to the core numbers is guaranteed. They also analyzed the convergence characteristics of the real-world networks and show nice trade-offs for time and quality of the solutions.

We generalize Lu et al.'s work for *any* nucleus decomposition, including k-truss. We show that convergence is guaranteed for all nucleus decompositions and prove the first upper bounds for the number of iterations. Our framework of algorithms locally compute *any* nucleus decomposition. We propose that iteratively computing h-indices of vertices/edges/r-cliques based on their degrees/triangle/s-clique counts converges in the core/truss/nucleus numbers ($r < s$). Local formulation also enables the pleasingly parallel computation, which is not possible in the peeling process.

(in seconds)	Peeling	Local alg.	Speedup
`twitter`	131	73	**1.8**
`web-NotreDame`	261	33	**7.9**
`wikipedia-200611`	3034	820	**3.7**

Table 1: **Runtime comparison between the peeling algorithm and the local algorithm we introduce for (3, 4) nucleus decomposition. Our local algorithm enables the pleasingly parallel computation and obtains up to 7.9x speedup over peeling with 24 threads.**

In addition, intermediate values provide an approximation to the exact nucleus decomposition to trade-off between runtime and quality. Note that this is also not possible in the peeling process, because no intermediate solution can provide an overall approximation to the exact solution. For instance, the densest regions are not revealed until the end.

1.2 Contributions

Our contributions can be summarized as follows:

- **Generalization for nucleus decomposition:** We generalize the iterative h-index computation idea [20] for *any* nucleus decomposition by using only the local information. Our approach is based on iteratively computing the h-indices on the degrees of vertices, triangle counts of edges, and s-clique counts of r-cliques ($r < s$) until convergence. We prove that the iterative computation by the h-indices guarantees the convergence to the core, truss, and nucleus decompositions.

- **Upper bounds for convergence:** We prove an upper bound for the number of iterations needed for convergence. We define the concept of *degree levels* which models the worst case for convergence. Our bounds are generic for *any* nucleus decomposition and much tighter than the obvious bounds that rely on the number of vertices/edge/triangles.

- **Framework of parallel local algorithms:** We introduce a framework of efficient algorithms that only use local information to compute *any* nucleus decomposition. Our algorithms are pleasingly parallel thanks to the local computation and have been implemented in OpenMP for the shared-memory architectures.

- **Evaluation on real-world networks:** We evaluate our algorithms and implementation on various types of real-world networks. With 24 threads, we obtain up to 4.04x and 7.98x speedups for k-truss and $(3, 4)$ nucleus decompositions, respectively. Table 1 gives representative runtime results for the $(3, 4)$ nucleus decomposition. We also obtain close approximations to the exact nucleus decompositions within a few iterations. Figure 1 presents the convergence rates by the number of iterations for k-truss and $(3, 4)$ nucleus decompositions. For most graphs, our algorithms achieve close solutions to the exact decompositions in $\sim$10 iterations. This enables the exploration of trade-offs between time and quality, which are essential for real-world applications.

Table 2: **Notations**

Symbol	Description
G	graph
R, S	r-clique, s-clique
$\mathcal{R}(G)$	set of r-cliques in graph G
$\mathcal{S}(G)$	set of s-cliques in graph G
$\mathcal{C}(G)$	$\mathcal{R}(G) \cup \mathcal{S}(G)$
$d_s\|_G(R)$ ($d_s(R)$)	# s-cliques that contains R in graph G (it is actually the S-degree of R)
$\delta_{r,s}(G)$	minimum S-degree of an r-clique in graph G
$\mathcal{N}_s(R)$	set of (neighbor) R's s.t. $\exists$ an s-clique $S \supseteq (R \cup R')$
$\kappa_s(R)$	largest k s.t. R is contained in a k-(r,s) nucleus
$\mathcal{H}(K)$	largest h s.t. at least h numbers in set K are $\geq h$
$\mathcal{U}$	update operator, $(\tau : \mathcal{R} \to \mathbb{N}) \to (\mathcal{U}\tau : \mathcal{R} \to \mathbb{N})$

2. BACKGROUND

In this section, we present several definitions to introduce our algorithms. We work on an undirected simple graph $G = (V, E)$ where V is the set of vertices and E is the set of edges.

2.1 Nucleus Decomposition

We start by giving the k-(r,s) nucleus subgraph definition that was proposed in [28].

DEFINITION 1. *Let $r < s$ be positive integers.*

- *$\mathcal{R}(G)$and $\mathcal{S}(G)$ are the set of r-cliques and s-cliques in G, respectively, and $\mathcal{C}(G) = \mathcal{R}(G) \cup \mathcal{S}(G)$ (We may use $\mathcal{C}$, $\mathcal{R}$, and $\mathcal{S}$ when G is obvious).*
- *In graph G, the number of $S \in \mathcal{S}(G)$ containing the r-clique $R \in \mathcal{R}(G)$ is the* ***$\mathcal{S}$-degree*** *of R, denoted as $\boldsymbol{d_s|_G(R)}$. (Abusing the notation, we may also use $d_s(R)$ when G is obvious in the context)*
- *In graph G, minimum $\mathcal{S}$-degree of an r-clique $R \in \mathcal{R}(G)$ is denoted as $\boldsymbol{\delta_{r,s}(G)}$.*
- *Two r-cliques R, R' are* ***$\mathcal{S}$-connected*** *if there exists a sequence $R = R_1, R_2, \ldots, R_k = R'$ in $\mathcal{R}$ such that for each i, some $S \in \mathcal{S}$ contains $R_i \cup R_{i+1}$.*
- ***$\mathcal{N}_s(R)$*** *denotes the set of neighbor r-cliques of the r-clique R such that $R' \in \mathcal{N}_s(R)$ if $\exists$ an s-clique S s.t. $S \ni R$ and $S \ni R'$.*

DEFINITION 2. *Let k, r, and s be positive integers such that $r < s$. A* ***k-(r,s) nucleus*** *is a subgraph G' which contains the edges in the maximal union $\mathcal{S}$ of s-cliques such that:*

- *The $\mathcal{S}$-degree of any r-clique $R \in \mathcal{R}(G')$ is at least k.*
- *Any r-clique pair $R, R' \in \mathcal{R}(G')$ are $\mathcal{S}$-connected.*

k-(r,s) nucleus definition is inspired by k-core and and k-truss. For $r = 1, s = 2$, k-$(1,2)$ nucleus is a maximal (induced) connected subgraph with minimum vertex degree k. This is exactly the k-core. Setting $r = 2, s = 3$ gives maximal subgraphs where every edge participates in at least k triangles, and edges are triangle-connected. This is essentially the definition of k-truss. We also define the κ_s index of r-clique analogous to the core numbers of vertices and truss numbers of edges.

Algorithm 1: PEELING(G, r, s)

```
Input: G: the graph, r, s: natural numbers
Output: κ_s(·): array of κ_s indices for r-cliques
Enumerate all r-cliques in G
For every r-clique R, set d_s(R) (S-degrees)
Mark every r-clique as unprocessed
for each unprocessed r-clique R with minimum d_s(R)
 do
    κ_s(R) = d_s(R)
    Find set S of s-cliques containing R
    for each C ∈ S do
        if any R ⊂ C is processed then continue
        for each r-clique R' ⊂ C, R' ≠ R do
            if d_s(R') > d_s(R) then d_s(R') = d_s(R') − 1
    Mark R as processed
return array κ_s(·)
```

DEFINITION 3. *For any r-clique $R \in \mathcal{R}(G)$, the* ***κ_s-index*** *of R, denoted as $\kappa_s(R)$, is the largest k value such that R is contained in a k-(r,s) nucleus.*

We use the notion of k-(r,s) nucleus and κ_s-index to introduce our generic theorems and algorithms for *any* r, s values. The set of k-(r,s) nuclei is found by the peeling algorithm [28] (outlined in Algorithm 1 for completeness). It is a generalization of the k-core and k-truss decomposition algorithms, and finds the κ_s indices of r-cliques in non-decreasing order.

The following lemma is quite standard in the k-core literature and we prove the analogue for k-(r,s) nucleus. It is a convenient characterization of the κ_s indices.

LEMMA 1. *For all $R \in \mathcal{R}(G)$, $\kappa_s(R) = max_{\mathcal{R}(G') \ni R} \delta_{r,s}(G')$, where $G' \subseteq G$.*

PROOF. Let T be the $\kappa_s(R)$-(r,s) nucleus containing R. By definition, $\delta_{r,s}(T) = \kappa_s(R)$, so $\max_{G'} \delta_{r,s}(G') \geq \kappa_s(R)$. For contradiction's sake, suppose there existed some subgraph $T' \ni R$ such that $\delta_{r,s}(T') > \kappa_s(R)$ (Wlog, we can assume T' to be connected; otherwise, we denote T' to be the component containing R). There must exist some maximal connected $T'' \supseteq T'$ that is a $\delta_{r,s}(T')$-nucleus. This would imply that $\kappa_s(R) \geq \delta_{r,s}(T') > \kappa_s(R)$, a contradiction. □

2.2 h-index computation

The main idea in our work is the iterative h-index computation on the S-degrees of r-cliques. h-index metric is introduced to measure the impact and productivity of researchers by the citation counts [14]. A researcher has h-index if she has at least h papers and each paper is cited at least h times such that there is no $h' > h$ that satisfies these conditions. We define the function $\mathcal{H}$ to compute the h-index as follows:

DEFINITION 4. *Given a set K of real numbers, $\mathcal{H}(K)$ is the largest $h \in \mathbb{N}$ such that at least h elements of K are at least h.*

Core numbers are related to h-indices. For instance, core number of a vertex can be defined as the largest k such that it has at least k neighbors whose core numbers are also at least k. In the following section, we will formalize this observation, and build on it to design algorithms to compute not only core decompositions but also truss and generalized nucleus decompositions.

3. GOING FROM h-INDEX TO κ_s-INDEX

Our main mathematical contribution is two-fold. First, we introduce a generic formulation to compute the k-(r, s) nucleus by an iterated h-index computation on r-cliques. Secondly, we prove convergence bounds for the number of iterations.

We define the *update operator* $\mathcal{U}$. This takes a function $\tau : \mathcal{R} \to \mathbb{N}$ and returns another function $\mathcal{U}\tau : \mathcal{R} \to \mathbb{N}$, where $\mathcal{R}$ is the set of r-cliques in the graph.

DEFINITION 5. *The update $\mathcal{U}$ is applied on the r-cliques in a graph G such that for each r-clique $R \in \mathcal{R}(G)$:*

1. *For each s-clique $S \ni R$, set $\rho(S, R) = \min_{R' \in S, R' \neq R} \tau(R')$.*
2. *Set $\mathcal{U}\tau(R) = \mathcal{H}(\{\rho(S, R)\}_{S \supset R})$.*

Observe that $\mathcal{U}\tau$ can be computed in parallel over all r-cliques in $\mathcal{R}(G)$. It is convenient to think of the $\mathcal{S}$-degrees (d_s) and κ_s indices as functions $\mathcal{R} \to \mathbb{N}$. We initialize $\tau_0 = d_s$, and set $\tau_{t+1} = \mathcal{U}\tau_t$.

The results of Lu et al. [20] prove that, for the k-core case ($r = 1, s = 2$), there is a sufficiently large t such that $\tau_t = \kappa_2$ (core number). We generalize this result for *any* nucleus decomposition. Moreover, we prove the first convergence bounds for $\mathcal{U}$.

The core idea of [20] is to prove that the $\tau_t(\cdot)$ values never increase (monotonicity) and are always lower bounded by core numbers. We generalize their proof for nuclei.

THEOREM 1. *For all t and all r-cliques R:*

- *(Monotonicity) $\tau_{t+1}(R) \leq \tau_t(R)$.*
- *(Lower bound) $\tau_t(R) \geq \kappa_s(R)$.*

PROOF. (Monotonicity) We prove by induction on t. Consider the base case $t = 0$. Note that for all R, $\tau_1(R) = \mathcal{U}d_s(R) \leq d_s(R)$. This is because in Step 2, the $\mathcal{H}$ operator acts on a set of $d_s(R)$, and this is largest possible value it can return. Now for induction. (Assume the property is true up to t.) Fix an r-clique R, and s-clique $S \supset R$. For $\tau_t(R)$, one computes the value $\rho(S, R) = \min_{R' \subset S, R' \neq R} \tau_{t-1}(R)$. By the induction hypothesis, the values $\rho(S, R)$ computed for τ_{t+1} is at most the value computed for τ_t. Note that the $\mathcal{H}$ operator is monotone; if one decreases values in a set K, then $\mathcal{H}(K)$ cannot increase. Since the ρ values cannot increase, $\tau_{t+1}(R) \leq \tau_t(R)$.

(Lower bound) We will prove that for any $G' \subseteq G$, $\mathcal{R}(G') \ni R$, $\tau_t(R) \geq \delta_{r,s}(G')$. Lemma 1 completes the proof.

We prove the above by induction on t. For the base case, $\tau_0(R) = d_s|_G(R) \geq d_s|_{G'}(R) \geq \delta_{r,s}(G')$. Now for induction. By the induction hypothesis, $\forall\ R \in \mathcal{R}(G'), \tau_t(R) \geq \delta_{r,s}(G')$. Consider the computation of $\tau_{t+1}(R)$, and the values $\rho(S, R)$ computed in Step 1. For every s-clique S, note that $\rho(S, R) = \min_{R' \subset S, R' \neq R} \tau_t(R')$. By the induction hypothesis, this is at least $\delta_{r,s}(G')$. By definition of $\delta_{r,s}(G')$, $d_s|_{G'}(R) \geq \delta_{r,s}(G')$. Thus, in Step 2, the $\mathcal{H}$ operator returns at least $\delta_{r,s}(G')$. □

Note that this is an intermediate result and we will present our final result in Lemma 2 at the end.

3.1 Convergence bounds by the degree levels

A trivial upper bound for convergence is the number of r-cliques in the graph, $|\mathcal{R}(G)|$, because after n iterations n r-cliques with the lowest κ_s indices will converge. We present a tighter bound for convergence. Our main insight is to define the *degree levels* of r-cliques, and relate these to the convergence of τ_t to κ_s. We prove that the κ_s indices in the ith level converge within i iterations of the update operation. This gives quantitative bounds on the convergence.

DEFINITION 6. *For a graph G,*

- *$S \in \mathcal{C}(G)$ if and only if $R \in \mathcal{C}(G), \forall\ R \subset S$.*
- *If R is removed from $\mathcal{C}(G)$, all $S \supset R$ are also removed from $\mathcal{C}(G)$.*
- **Degree levels** *are defined recursively as follows. The i-th level is the set L_i.*
  - *L_0 is the set of r-cliques that has the minimum $\mathcal{S}$-degree in $\mathcal{C}$.*
  - *L_i is the set of r-cliques that has the minimum $\mathcal{S}$-degree in $\mathcal{C} \setminus \bigcup_{j<i} L_j$.*

We first prove the κ_s indices cannot decrease as the level increases. The following proof is closely related to the fact the minimum degree removal algorithm (peeling) finds all cores/nuclei.

THEOREM 2. *Let $i \leq j$. For any $R_i \in L_i$ and $R_j \in L_j$, $\kappa_s(R_i) \leq \kappa_s(R_j)$.*

PROOF. Let $L' = \bigcup_{r \geq i} L_r$, the union of all levels i and above, and G' is the graph such that $L' = \mathcal{R}(G')$. By definition of the levels, $d_s|_{G'}(R_i) = \delta_{r,s}(G')$ and $d_s|_{G'}(R_j) \geq d_s|_{G'}(R_i)$. There exists some $\kappa_s(R_i)$-nucleus T containing R_i. We split into two cases.

Case 1: $\mathcal{R}(T) \subseteq L'$. Thus, $\kappa_s(R_i) = \delta_{r,s}(T) \leq \delta_{r,s}(G') = d_s|_{G'}(R_i)$. Note that $\kappa_s(R_j) = \min_{P \ni R_j} \delta_{r,s}(P)$, so $\kappa_s(R_j) \geq \delta_{r,s}(G')$. Hence, $\kappa_s(R_i) \leq \kappa_s(R_j)$.

Case 2: $\mathcal{R}(T) \setminus L' \neq \emptyset$. Thus, there exists some r-clique $R' \in \mathcal{R}(T) \cap L_b$, where $b < i$. Choose the R' that minimizes this value of b. Since T is a $\kappa_s(R_i)$-nucleus, $d_s|_T(R') \geq \kappa_s(R_i)$. Consider $M = \bigcup_{r \geq b} L_r$. Note that $\mathcal{R}(T) \subseteq M$, since we chose R' to minimize b. Let Q is the graph such that $M = \mathcal{R}(Q)$. We have $d_s|_Q(R') \geq d_s|_T(R') \geq \kappa_s(R_i)$. Since $R' \in L_b$, $d_s|_Q(R') = \delta_{r,s}(Q)$. Since $j > b$ and $R_j \in M$, $\kappa_s(R_j) \geq \delta_{r,s}(Q)$. Combining the above, we deduce $\kappa_s(R_i) \leq \kappa_s(R_j)$. □

The main convergence theorem is the following. As explained earlier, it shows that the ith level converges within i iterations.

THEOREM 3. *Fix any level L_i. For all $t \geq i$ and $R \in L_i$, $\tau_t(R) = \kappa_s(R)$.*

PROOF. We prove by induction on i. For the base case $i = 0$: note that for any R of minimum $\mathcal{S}$-degree in G, $\kappa_s(R) = d_s|_G(R) = \tau_0(R)$. Now for induction. Assume the theorem is true up to level i. Thus, for $t \geq i$ and $\forall\ R \in \bigcup_{j \leq i} L_j$, $\tau_t(R) = \kappa_s(R)$. Select arbitrary $R_a \in L_{i+1}$, and set $L' = \bigcup_{j \geq i+1} L_j$.

We partition the s-cliques containing R_a into the "low" set $\mathcal{S}_\ell$ and "high" set $\mathcal{S}_h$. s-cliques in $\mathcal{S}_\ell$ contain some r-clique outside L', and those in $\mathcal{S}_h$ are contained in L'. For every s-clique $S \in \mathcal{S}_\ell$, there is some $R_b \in S$ such that $R_b \in L_k$ for $k \leq i$. By the induction hypothesis, $\tau_t(R_b) = \kappa_s(R_b)$. By

Theorem 2 applied to $R_b \in L_k$ and $R_a \in L_{i+1}$, $\kappa_s(R_b) \leq \kappa_s(R_a)$.

Focus on the computation of $\tau_{t+1}(R_a)$, which involves first computing $\rho(S, R_a)$ in Step 1 of Definition 5. For every $S \in \mathcal{S}_\ell$, by the previous argument, there is some r-clique $R_b \in S$, $R_b \neq R_a$, such that $\tau_t(R_b) \leq \kappa_s(R_a)$. Thus, $\forall\, S \in \mathcal{S}_\ell$, $\rho(S, R_a) \leq \kappa_s(R_a)$. This crucially uses the *min* in the setting of $\rho(S, R_a)$, and is a key insight into the generalization of iterated $\mathcal{H}$-indices for any nucleus decomposition.

The number of edges in $\mathcal{S}_h$ is exactly $d_s|_{G'}(R_a) = \delta_{r,s}(G')$. Applying Lemma 1 to $R_a \in L'$, we deduce $\kappa_s(R_a) \geq d_s|_{G'}(R_a)$. All in all, for all $S \in \mathcal{S}_\ell$, $\rho(S, R_a)$ is at most $\kappa_s(R_a)$. On the other hand, there are at most $\kappa_s(R_a)$ s-cliques in $\mathcal{S}_h$. The application of the $\mathcal{H}$ function in Step 2 yields $\tau_{t+1}(R_a) \leq \kappa_s(R_a)$. But the lower bound of Theorem 1 asserts $\tau_{t+1}(R_a) \geq \kappa_s(R_a)$, and hence, these are equal. This completes the induction. □

We have the following lemma to show that convergence is guaranteed in a finite number of iterations.

LEMMA 2. *Given a graph G let l be the maximum i, such that $L_l \neq \emptyset$ and $\tau_l(R) \geq \kappa_s(R)$ for all r-cliques (e.g., $\tau_0 = d_s$) and set $\tau_{t+1} = \mathcal{U}\tau_t$. For some $t \leq l$, $\tau_t(R) = \kappa_s(R)$, for all r-cliques.*

4. LOCAL ALGORITHMS

We introduce generalized local algorithms to find the κ_s indices of r-cliques for any (r, s) nucleus decomposition. For each r-clique, we iteratively apply h-index computation on some attributes in their neighborhood. Our local algorithms are pleasingly parallel thanks to the independent nature of the h-index computations. We also offer to explore time and quality trade-offs by using the iterative nature. We first present the deterministic synchronous algorithm which does not depend on the order of processing the r-cliques. It implements the $\mathcal{U}$ operator in Definition 5. Then we adapt our algorithm to work in an asynchronous manner that converges faster and uses less space. For those familiar with linear algebra, the synchronous and asynchronous algorithms are analogous to Jacobi and Gauss-Seidel iterations for iterative solvers. At the end, we discuss some heuristics and key implementation details for shared-memory parallelism in OpenMP.

4.1 Synchronous Nucleus Decomposition (SND)

We use the update operator $\mathcal{U}$ to compute the k-(r, s) nuclei of a graph G in a synchronous way. Algorithm 2 (SND) implements the Definition 5 for functions $\tau_0 = d_s$ and $\tau_{t+1} = \mathcal{U}\tau_t$ to find the κ_s indices of r-cliques in graph G.

SND algorithm proceeds in iterations until no further updates occur for any τ index, which means all the τ indices converged to κ_s. Computation is performed synchronously on all the r-cliques and at each iteration i, τ_i indices are found for all r-cliques. We declare two arrays, $\tau(\cdot)$ and $\tau^p(\cdot)$, to store the indices being computed and the indices that were computed in the previous iteration, respectively (Lines 1 and 4). $\tau(\cdot)$ are initialized to the $\mathcal{S}$-degrees of the r-cliques since $\tau_0 = d_s$ (Line 2). At each iteration, newly computed $\tau(\cdot)$ indices are backed up in $\tau^p(\cdot)$ (Line 7), and the new $\tau(\cdot)$ indices are computed. During the iterative process, convergence is checked by the flag $\mathcal{F}$ (Line 5), which is initially set to TRUE (Line 3) and stays TRUE as long as there is an update on a τ index (Lines 6, 13, and 14).

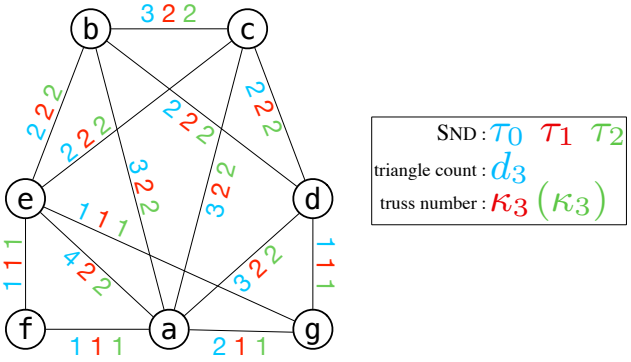

Figure 2: **SND (Algorithm 2) for the k-truss decomposition ($r = 2, s = 3$). We find the κ_3 indices. Triangle counts of all the edges are computed (d_3) and set as their τ_0 values (blue). For each edge, we first compute τ_1 indices (red) based on the τ_0 indices. The *ae* edge, for instance, has four triangles and for each of those we find the neighbor with minimum τ_0 index and compute the h-index. So, $\tau_1(ae) = \mathcal{H}\{(min(\tau_0(eb), \tau_0(ab)), min(\tau_0(ec), \tau_0(ac)), min(\tau_0(eg), \tau_0(ag)), min(\tau_0(ef), \tau_0(af))\} = \mathcal{H}\{2, 2, 1, 1\} = 2$. No updates happen in the second iteration (green), so convergence is obtained in a single iteration.**

Computation of the new $\tau(\cdot)$ indices for each r-clique can be performed in parallel with no communication or synchronization between threads (Lines 8 to 15). For each r-clique R, we apply the two step process in the Definition 5. First, for each s-clique S that contains R, we compute the ρ values which is the minimum τ^p index of an r-clique $R' \in S$ ($R' \neq R$) and collect them in a set L (Lines 10 to 12). Then, we compute the h-index of the set L and assign it as the new τ index of the r-clique (Line 15). We also check if there is an update on the τ index and make sure to continue computation if so (Line 13 and 14). Once the τ indices converge, we assign them to κ_s indices and complete the algorithm (Lines 16 and 17)

Figure 2 illustrates the SND algorithm for k-truss decomposition ($r = 2, s = 3$) on a toy graph, where the involvements of edges (2-cliques) in triangles (3-cliques) are examined. Triangle counts of all the edges (d_3) are computed and set as their τ_0 values (in blue). For each edge, first we compute τ_1 indices (in red) based on the τ_0 indices (Line 5 to 15). For instance, the **ae** edge has four triangles and for each of those we find the neighbor with minimum τ_0 index (Line 10 to 12). So, set $L = \{(min(\tau_0(\mathbf{eb}), \tau_0(\mathbf{ab})), min(\tau_0(\mathbf{ec}), \tau_0(\mathbf{ac})), min(\tau_0(\mathbf{eg}), \tau_0(\mathbf{ag})), min(\tau_0(\mathbf{ef}), \tau_0(\mathbf{af}))\} = \{2, 2, 1, 1\}$ and $\tau_1(\mathbf{ae}) = \mathcal{H}(L) = 2$ (Line 15). Since the τ index is updated, we set flag $\mathcal{F}$ TRUE to continue iterations. After computing the τ_2 indices (green), we observe that there is no update, i.e., $\tau_2(\mathbf{e}) = \tau_1(\mathbf{e})$ for all edges, thus completing the algorithm. So, one iteration is enough for the convergence and we have $\kappa_3 = \tau_1$ for all edges.

4.2 Asynchronous Nucleus Decomposition (AND)

In the SND algorithm, updates on the τ indices are synchronous and all the r-cliques are processed by using the same snapshot of τ indices. However, when an r-clique R is being processed in iteration i, a neighbor r-clique R' that

Algorithm 2: SynchronousNucDec(Snd) (G, r, s)

Input: G: the graph, r, s: natural numbers
Output: $\kappa_s(\cdot)$: array of κ_s indices for r-cliques

1 $\tau(\cdot) \leftarrow$ indices $\forall\ R \in \mathcal{R}(G)$ `// for current iter.`
2 $\tau(R) \leftarrow d_s(R)\ \forall\ R \in \mathcal{R}(G)$ `// set to the` $\mathcal{S}$`-degrees`
3 $\mathcal{F} \leftarrow$ TRUE `// stays` TRUE `if any` $\tau(R)$ `is updated`
4 $\tau^p(\cdot) \leftarrow$ indices $\forall\ R \in \mathcal{R}(G)$ `// for previous iter.`
5 **while** $\mathcal{F}$ **do**
6 $\quad\mathcal{F} \leftarrow$ FALSE
7 $\quad\tau^p(R) \leftarrow \tau(R)\ \forall\ R \in \mathcal{R}(G)$
8 $\quad$**for each** $R \in \mathcal{R}(G)$ *in parallel* **do**
9 $\quad\quad L \leftarrow$ empty set
10 $\quad\quad$**for each** s-clique $S \ni R$ **do**
11 $\quad\quad\quad\rho \leftarrow \min_{R' \in S \wedge R' \neq R} \tau^p(R')$
12 $\quad\quad\quad L$. add (ρ)
13 $\quad\quad$**if** $\tau^p(R) \neq \mathcal{H}(L)$ **then**
14 $\quad\quad\quad\mathcal{F} \leftarrow$ TRUE
15 $\quad\quad\tau(R) \leftarrow \mathcal{H}(L)$
16 $\kappa_s(\cdot) \leftarrow \tau(\cdot)$
17 **return** array $\kappa_s(\cdot)$

Algorithm 3: AsynchronousNucDec(And) (G, r, s)

Input: G: the graph, r, s: natural numbers
Output: $\kappa_s(\cdot)$: array of κ_s indices for r-cliques

1 $\tau(\cdot) \leftarrow$ indices $\forall\ R \in \mathcal{R}(G)$ `// for current iter.`
2 $\tau(R) \leftarrow d_s(R)\ \forall\ R \in \mathcal{R}(G)$ `// set to the` $\mathcal{S}$`-degrees`
3 $\mathcal{F} \leftarrow$ TRUE `// stays` TRUE `if any` $\tau(R)$ `is updated`
4 $\mathtt{c}(R) \leftarrow$ TRUE $\forall\ R \in \mathcal{R}(G)$
5 **while** $\mathcal{F}$ **do**
6 $\quad\mathcal{F} \leftarrow$ FALSE
7 $\quad$**for each** $R \in \mathcal{R}(G)$ *in parallel* **do**
8 $\quad\quad$**if** $\mathtt{c}(R)$ *is* FALSE **then** continue
9 $\quad\quad L \leftarrow$ empty set
10 $\quad\quad$**for each** s-clique $S \ni R$ **do**
11 $\quad\quad\quad\rho \leftarrow \min_{R' \in S \wedge R' \neq R} \tau(R')$
12 $\quad\quad\quad L$. add (ρ)
13 $\quad\quad$**if** $\tau(R) \neq \mathcal{H}(L)$ **then**
14 $\quad\quad\quad\mathcal{F} \leftarrow$ TRUE
15 $\quad\quad\quad$**for each** $R' \in \mathcal{N}_s(R)$ **do**
16 $\quad\quad\quad\quad$**if** $\mathcal{H}(L) < \tau(R') \leq \tau(R)$ **then**
17 $\quad\quad\quad\quad\mathtt{c}(R') \leftarrow$ TRUE
18 $\quad\quad\quad\tau(R) \leftarrow \mathcal{H}(L)$
19 $\quad\quad\quad\mathtt{c}(R) \leftarrow$ FALSE
20 $\kappa_s(\cdot) \leftarrow \tau(\cdot)$
21 **return** array $\kappa_s(\cdot)$

participates in an s-clique with R might have already completed its computation in that iteration and updated its τ index. By Theorem 1, we know that the τ index can only decrease as the algorithm proceeds. Lower $\tau(R')$ indices in set L might decrease $\mathcal{H}(L)$, and it can help $\tau(R)$ to converge faster. So, it is better to use the *up-to-date* τ indices for faster convergence. In addition, there would be no need to store the τ indices computed in the previous iteration, saving $r|\mathcal{R}(G)|$ space.

We introduce the And algorithm (Algorithm 3) to leverage the up-to-date τ indices for faster convergence (Orange lines can be ignored for now). At each iteration, we propose to use the latest available information in the neighborhood of an r-clique. Removing the green lines in the Snd algorithm and inserting the blue lines in the And algorithm are

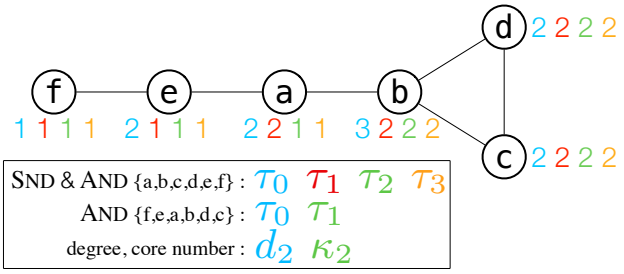

Figure 3: **Snd (Algorithm 2) and And (Algorithm 3, no orange lines) for the k-core decomposition ($r = 1, s = 2$). We find the κ_2 indices (core numbers) of vertices (edge is 2-clique). τ_0 indices are initialized to the degrees (d_2 in blue). Snd algorithm uses the τ_{i-1} indices to compute the τ_i indices and converges in two iterations (τ_1 in red, τ_2 in green, τ_3 in yellow). Same happens when we use the And algorithm (without orange lines) and follow the $\{a,b,c,d,e,f\}$ order to process the vertices. If we choose $\{f,e,a,b,c,d\}$ order, which is actually a non-decreasing order on κ_2 indices, And converges in a single iteration.**

sufficient to switch from synchronous to asynchronous computation. We do not need to use the $\tau_p(\cdot)$ to back up the indices in the previous iteration anymore, so lines 4 and 7 in Algorithm 2 are removed. Computation is done on the latest τ indices, so we adjust the lines 11 and 13 (in Algorithm 2 and 3) accordingly, to use the up-to-date τ indices.

In the same iteration, each r-clique can have a different view of the $\tau(\cdot)$ and updates are done *asynchronously* in an arbitrary order. Number of iterations for convergence depends on the processing order of the r-cliques, which is used in line 7 of Algorithm 3. For the sequential computation, we have the following theorem regarding the best ordering that results in the quickest convergence.

Theorem 4. *For a graph G, if the r-cliques are processed in the non-decreasing order of their κ_s indices in the* And *algorithm, convergence is obtained in a single iteration.*

Proof. Say that we process the r-cliques in the non-decreasing order of κ_s indices and for the sake of contradiction, assume that there exists an r-clique R s.t. $\tau_1(R) > \kappa_s(R)$, i.e., it does not converge in a single iteration. Let R_0 is the first r-clique being processed, so we have $\kappa_s(R) \geq \kappa_s(R_0)$ since we follow the non-increasing order.

(1) If $\kappa_s(R) = \kappa_s(R_0)$: Say $\kappa_s(R) = n$. $d_s(R) = \tau_0(R) = m \geq n$ by Theorem 1. So initially there are m s-cliques that contains R. We also know that for each $R' \in \mathcal{R}(G)$, $\tau_0(R') \geq \kappa_s(R') \geq n$ since n is the smallest κ_s index, by our assumption. Thus, for each s-clique $S \ni R$, ρ value is at least n and there are m such s-cliques, which makes $\tau_1(R) = n$. But, our initial assumption is $\tau_1(R) > \kappa_s(R) = n$, contradiction.

(2) If $\kappa_s(R) > \kappa_s(R_0)$: Say $\kappa_s(R_0) = n$ and $\kappa_s(R) \geq n+1$. We have $\tau_0(R) \geq \tau_1(R) > n + 1$, by our assumption, when we start processing R. Since R is not the first r-clique in the order, we have a set P of r-cliques that are processed before R where $|P| > 0$. Thus, we have $d_s(R) > n + |P| + 1$

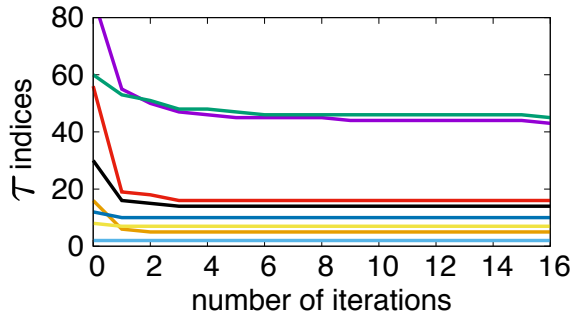

Figure 4: **Changes in τ indices of randomly selected edges in `facebook` graph during the k-truss decomposition. There are wide plateaus during the convergence, especially at the end, where τ indices stay constant.**

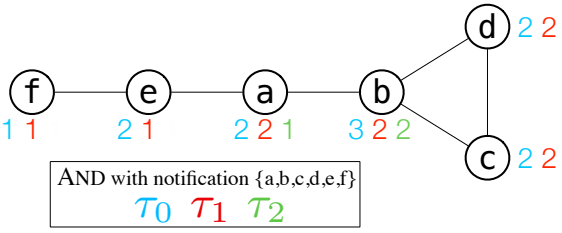

Figure 5: **k-core decomposition ($r = 1, s = 2$) by AND (Algorithm 3) that uses the notification mechanism. After the first iteration, the only active vertex is a and recomputing updates its τ index and also notifies vertex b. Recomputing τ of vertex b does not yield any change and all the vertices become idle, thus convergence is obtained. 8 τ computations performed in 3 iterations by AND with notification mechanism, while 24 τ computations done in 4 iterations with if no notifications used (Figure 3).**

at most, when $\forall\ R' \in P$, $\kappa_s(R') < \kappa_s(R)$ and there exists an s-clique $S \in G$ s.t. $R \cup R' \subseteq S$. If we exclude the s-cliques that contains an r-clique from P, we have more than $n+1$ s-cliques where each has a ρ value of at least $n+1$, since all the r-cliques R'' in those s-cliques has $d_s(R'') > n$. This implies that $\tau_1(R) = n+1$, which contradicts with our initial assumption that $\tau_1(R) > \kappa_s(R) \geq n+1$. □

The worst case for AND happens when all the r-cliques see the τ indices of their neighbors that are computed in the previous iteration, and it is exactly the SND algorithm.

Figure 3 illustrates the SND algorithm and AND algorithm with two different orderings on the k-core case ($r = 1, s = 2$). Focus is on vertices (1-cliques) and their edge (2-clique) counts (degrees). We first do the SND. Vertex degrees are set as τ_0 indices (blue). For each vertex u we compute the $\tau_1(u) = \mathcal{H}(\{\tau_0(v) : v \in \mathcal{N}_2(u)\}$ (red), i.e., h-index of its neighbors' degrees. For instance, vertex **a** has two neighbors, **e** and **b**, with degrees 2 and 3. Since $\mathcal{H}(\{2,3\}) = 2$, we get $\tau_1(\mathbf{a}) = 2$. For vertex **b**, we get $\tau_1(\mathbf{b}) = \mathcal{H}(\{2,2,2\}) = 2$. Once we compute all the τ_1 indices, we observe an update (vertices **e** and **b**), thus move forward and compute the τ_2 indices, shown in green. We observe an update only for the vertex **a**; $\tau_2(\mathbf{a}) = \mathcal{H}(\{\tau_1(\mathbf{e}), \tau_1(\mathbf{b})\}) = \mathcal{H}(\{1,2\}) = 1$ and continue computation. For τ_3 indices (yellow), no update is observed which means that $\kappa_s = \tau_2$ for all vertices and SND converges in two iterations. Regarding the AND algorithm, say we choose to follow the non-decreasing order of κ_s indices; {**f,e,a,b,c,d**}. Computing the τ_1 indices on this order enables us to reach the convergence in a single iteration. For instance, $\tau_1(\mathbf{a}) = \mathcal{H}(\{\tau_1(\mathbf{e}), \tau_0(\mathbf{b})\}) = \mathcal{H}(\{1,3\}) = 1$. However, if we choose to process the vertices in the alphabetical order, {**a,b,c,d,e,f**}, we have $\tau_1(\mathbf{a}) = \mathcal{H}(\{\tau_0(\mathbf{e}), \tau_0(\mathbf{b})\}) = \mathcal{H}(\{2,3\}) = 2$, and need more iteration(s) to converge. Indeed, **a** is the only updated vertex. In the second iteration, we get $\tau_2(\mathbf{a}) = \mathcal{H}(\{\tau_1(\mathbf{e}), \tau_1(\mathbf{b})\}) = \mathcal{H}(\{1,2\}) = 1$, an update, thus continue iterating. Third iteration does not change the τ indices, so AND with {**a,b,c,d,e,f**} order converges in two iterations, same as the SND.

4.2.1 *Skipping the plateaus*

SND and AND algorithms converge when none of the r-cliques update their τ indices anymore. Consequently, computations continue to be performed for all the r-cliques even when only one update occurs and we also need an extra iteration to detect the convergence. Figure 4 shows the τ indices of randomly selected edges in the `facebook` graph during k-truss decomposition ($r = 2, s = 3$). There are plenty of wide plateaus where τ indices stay constant, which suggests that those computations are redundant. But how can we avoid them? For example, when $\tau(R)$ converges to $\kappa_s(R)$, no more computations are needed for R anymore. However, we cannot know if the $\tau(R)$ has converged or not by looking at the repeating τ indices or watching the plateaus, because an update can occur after maintaininig the same τ index for a number of iterations, creating a plateau. In order to efficiently detect the convergence and skip any plateaus during the computation, we introduce a notification mechanism where an r-clique is notified to recompute its τ index, if any of its neighbors has an update.

Orange lines in Algorithm 3 present the notification mechanism plugged in to the asynchronous computation. $\mathtt{c}(\cdot)$ array is declared in (Line 4) to track whether an $R \in \mathcal{R}(G)$ has updated its τ index or not. $\mathtt{c}(R) = \text{FALSE}$ means R did not update its τ index, it is an idle r-clique, and there is no need to recompute its τ value for that iteration (Line 8). A non-idle r-clique is called active. Thus, all $\mathtt{c}(\cdot)$ is set to TRUE at the beginning to initiate the computations for all r-cliques. Each r-clique marks itself idle at the end of an iteration (Line 19) and waits until an update happens in the τ index of a neighbor. When the $\tau(R)$ is updated, τ indices of *some* neighbor r-cliques in $\mathcal{N}_s(R)$ might be affected and they should be notified for the recomputation. For $R' \in \mathcal{N}_s(R)$, if $\tau(R') \leq \mathcal{H}(L)$ (new $\tau(R)$) then it is not affected since $\tau(R') \leq \tau(R)$ already in the previous iteration and no change can happen in the h-index computation. In addition, if $\tau(R') > \tau(R)$ before, then $\tau(R') > \mathcal{H}(L)$ (the new $\tau(R)$) as well and will not be affected. Therefore, we only need to notify the neighbors that have τ indices in $(\mathcal{H}(L), \tau(R)]$ interval (Line 15 to 17) In parallel computation, some neighbors might already be active when they are notified, and use the stale $\tau(R)$ at that time, thus missing the new update. But it is okay since the following iterations will handle it. Indeed, parallel asynchronous computation becomes the synchronous computation *in the worst case*.

For the same toy graph, we now apply AND algorithm with the notification mechanism in the **{a,b,c,d,e,f}** order.

Figure 5 illustrates the computation. Again, vertex degrees are set as τ_0 indices (blue) and we compute the $\tau_1(u) = \mathcal{H}(\{\tau_0(v) : v \in \mathcal{N}_2(u)\}$, i.e., h-index of its neighbors' degrees, (red) for each vertex u. No update happens for vertex **a** and no vertex is notified. $\tau(\mathbf{b})$ is updated as 2 and we check if any neighbors of **b** has a τ index > 2 and ≤ 3 (Line 16). No such vertex occurs and no notification happens. Vertices **c** and **d** do not update their τ indices. $\tau(\mathbf{e})$ is updated as 1 and since $\tau_0(\mathbf{e}) \geq \tau_1(\mathbf{a}) > \tau_1(\mathbf{e})$, vertex **a** is notified for recomputing its τ index. Vertex **f** does not change its τ index as well and all the vertices except **a** are idle now. In the second iteration, we process **a** and compute $\tau_2(\mathbf{a}) = \mathcal{H}\{\tau_1(\mathbf{e}), \tau_1(\mathbf{b})\} = \mathcal{H}\{1, 2\} = 1$. Update in τ index of **a** notifies vertex **b** since $\tau_1(\mathbf{a}) \geq \tau_1(\mathbf{b}) > \tau_2(\mathbf{a})$. However computing $\tau_2(\mathbf{b})$ gives no update and all the vertices become idle, meaning that the convergence is achieved. Overall, we do 8 τ computations by using the AND with notification mechanism and it takes 3 iterations to complete the process, whereas 4 iterations needed if no notification mechanism is used and 24 τ computations are necessary in total (Figure 3). So the notification mechanism is helpful to avoid redundant computations.

4.3 Heuristics and implementation

We introduce key implementation details for the shared-memory parallelism and important heuristics for efficient h-index computation. We used OpenMP [6] to utilize the shared-memory architectures. The loops, annotated as parallel in Algorithms 2 and 3, are shared among the threads, and each thread is responsible for its partition of r-cliques. No synchronization or atomic operation is needed. Default scheduling policy in OpenMP is *static* and it distributes the iterations of the loop to the threads in chunks, i.e., for two threads, one takes the first half and the other takes the second. Although this policy is useful for many applications, it does not work well for our algorithms. The notification mechanism to avoid the redundant computations can result in significant load imbalance between threads with static scheduling. If most of the idle r-cliques are assigned to a certain thread, then that thread quickly finishes its work, becomes idle and waits for the other threads until the end of the iteration. To prevent this, we embraced the dynamic scheduling where each thread is given a new workload once it is done. We set chunk size to 100 and observed no significant difference for other values. No thread stays idle this way, and the overall computation is parallelized more efficiently.

h-index computation for a list of numbers is traditionally done by sorting the numbers in the non-increasing order and checking the values starting from the head of the list to find the largest h value for which at least h items exist with at least h value. Main bottleneck in this routine is the sorting operation which takes $O(nlogn)$ time for n numbers. We introduce a linear time algorithm that uses a hashmap and does not include sorting to compute the h-index. h is initialized as zero and we iterate over the items in the list. At each step, we attempt to increase the current h value based on the inspected item. For the present h value in a step, we keep track of the number of items examined so far that have value equal to h. We use a hashmap to keep track of the number of items that has at least h value, and we ignore values smaller than h. This enables the computation of the h-index in linear time and provides a nice trade-off between time and space. In addition, for the non-initial iterations of the convergence process, we simply check the items if the current τ index can be preserved. Once we see at least τ items with index at least τ, no more checks needed and no update happens.

Table 3: **Important statistics about our dataset; number of vertices, edges, triangles and four-cliques (K_4).**

	$\|V\|$	$\|E\|$	$\|\triangle\|$	$\|K_4\|$
`as-skitter` (ASK)	1.7M	11.1M	28.8M	148.8M
`facebook` (FB)	4K	88.2K	1.6M	30.0M
`soc-LiveJournal` (SLJ)	4.8M	68.5M	285.7M	9.9B
`soc-orkut` (ORK)	2.9M	106.3M	524.6M	2.4B
`soc-sign-epinions` (SSE)	131.8K	711.2K	4.9M	58.6M
`soc-twitter-higgs` (HG)	456.6K	12.5M	83.0M	429.7M
`twitter` (TW)	81.3K	1.3M	13.1M	104.9M
`web-Google` (WGO)	916.4K	4.3M	13.4M	39.9M
`web-NotreDame` (WND)	325.7K	1.1M	8.9M	231.9M
`wikipedia-200611` (WIKI)	3.1M	37.0M	88.8M	162.9M

5. EXPERIMENTS

We evaluate our algorithms on three instances of the (r, s) nucleus decomposition: k-core (or $(1, 2)$), k-truss (or $(2, 3)$), and $(3, 4)$. We do not store the s-cliques during the computation for better scalability in terms of the memory space. Instead, we find the participations of the r-cliques in the s-cliques on-the-fly. More details about the time-space trade-offs are given in [28]. Our dataset comprises a diverse set of real-world networks from SNAP [19] and Network Repository [25], such as an internet topology network (`as-skitter`), online social networks (`facebook`, `soc-LiveJournal`, `soc-orkut`), trust network (`soc-sign-epinions`), follower-followee Twitter networks (`soc-twitter-higgs`, `twitter`), web networks (`web-Google`, `web-NotreDame`), and a network of wikipedia pages (`wikipedia-200611`). Number of vertices, edges, triangles and four-cliques in those graphs are given in Table 3.

All experiments are performed on a Linux operating system running on a machine with Intel Ivy Bridge processor at 2.4 GHz with 64 GB DDR3 1866 MHz memory. There are two sockets on the machine and each has 12 cores, making 24 cores in total. Algorithms are implemented in C++ and compiled using gcc 6.1.0 at the -O2 optimization level. We used OpenMP v4.5 for the shared-memory parallelization.

We first compare the SND (Algorithm 2) and AND (Algorithm 3) algorithms and check the number of iterations for convergence in sequential settings. Then, we investigate the characteristics of the convergence processes and look how the τ indices are getting closer to the κ_s indices at each iteration. We explore the trade-offs between the quality of τ indices (how close they are to κ_s) and runtime, and give some guidance on finding the sweet spots where good-enough τ indices can be obtained faster. Last, but not least, we compare the runtimes of our algorithms with respect to the widely accepted peeling technique, and discuss the scalability of our implementations with increasing number of threads.

	Algorithm/Bound	ASK	FB	SLJ	ORK	SSE	HG	TW	WGO	WND	WIKI
k-core	Degree Levels	1195	352	3479	5165	642	1713	961	384	665	2026
	SND	63	21	99	147	38	73	37	23	187	55
	AND	33	11	51	73	19	37	21	14	35	30
k-truss	Degree Levels	1605	859	5401	4031	1273	2215	1337	254	417	2824
	SND	118	33	86	207	52	101	52	20	187	562
	AND	58	19	44	103	28	53	29	11	15	410
$(3,4)$	Degree Levels	1734	1171	7426	3757	1584	2360	1641	157	138	1559
	SND	72	38	123	196	54	109	76	11	14	122
	AND	41	23	73	116	30	51	40	6	10	107

Table 4: **Number of iterations given by upper bound (Section 3.1), and SND and AND algorithms. Both algorithms converge in far fewer iterations than the theoretical upper bounds – on average SND converges in 5% of the bounds for all decompositions. AND algorithm converges in 50% less number of iterations than the SND for k-core and k-truss decompositions and 35% less for $(3,4)$ nucleus decomposition.**

5.1 AND converges faster than SND

Number of iterations for convergence can be decreased by the asynchronous algorithm AND, as described in 4.2. We compare SND (Algorithm 2) and AND (Algorithm 3) for three nucleus decompositions. All the runs are performed in sequential, and for AND we use the natural ordering of the r-cliques in datasets that is the order of vertices/edges/triangles given or computed based on the ids in the data files. Note that, we also checked AND with random r-clique orderings and did not observe significant differences. Apart from the algorithms, we also compute the number of degree levels (Definition 6) that we prove as an upper bound in Section 3.1.

Table 4 presents the results for each decomposition type, k-core, k-truss and $(3,4)$. Number of degree levels gives much tighter bounds than the obvious limits – number of r-cliques. We observe that both algorithms converge in far fewer iterations than our upper bounds – SND converges in 5% of the bounds given for all decompositions, on average. Regarding the comparison, AND algorithm converges in 50% fewer iterations than the SND for k-core and k-truss decompositions and 35% less iterations for $(3,4)$ nucleus decomposition. Overall, we see the benefit of asynchronous computation on all decompositions and we use AND algorithm for the remaining experiments.

5.2 Few iterations are enough for close results

We investigate the convergence rates of the three decompositions. At each iteration i we compute the similarity between the τ_i indices and the final κ_s indices of r-cliques. We have two metrics to measure the similarity: (1) Kendall-Tau metric with a minor adaptation, (2) Accuracy percentage of the converged τ indices, i.e., $|\{R' \in \mathcal{R}(G) : \tau(R') = \kappa_s(R')\}|/|\mathcal{R}(G)|$.

Kendall-Tau metric computes the ratio of pairwise agreements between two ranking lists and used to measure the k-core approximations [20, 13]. In our case, for $(x_1, x_2, ..., x_{|\mathcal{R}(G)|})$ and $(y_1, y_2, ..., y_{|\mathcal{R}(G)|})$ lists, if both $(x_i < x_j)$ and $(y_i < y_j)$ or both $(x_i > x_j)$ and $(y_i > y_j)$, we say i, j pair is concordant. If $(x_i < x_j)$ and $(y_i > y_j)$ or $(x_i > x_j)$ and $(y_i < y_j)$, i, j pair is discordant, and if $(x_i = x_j)$ or $(y_i = y_j)$, then the pair is neither concordant nor discordant. Kendall-Tau is defined as $(n^+ - n^-)/\#\ pairs$, where n^+ and n^- are the number of concordant and discordant pairs. It is in $[-1, 1]$ interval and larger the number, more similar the two lists. For our problem, we make the Kendall-Tau metric stricter by considering that (i, j) pair is also concordant if both $(x_i = x_j)$ and $(y_i = y_j)$. For any other case, we say the pair is discordant. The reason behind our adaptation is that if two r-cliques have the same κ_s index, then they are in the same nucleus, which is critical to locate dense nuclei and relate them to each other in a hierarchical way. Note that, this metric becomes 1 at convergence and it is not necessarily monotonic with the increasing number of iterations in our algorithms, because when an r-clique decreases its τ index, many concordant pairs can be discordant temporarily for a number of iterations (especially the ones with equal indices before decrease).

Figure 6 presents the Kendall-Tau similarity scores between τ_i and κ_s indices for each iteration i, starting from 0 which shows the $\mathcal{S}$-degrees of r-cliques. We present the results for k-core and k-truss computations in Figure 6 and $(3,4)$ results are also given in Figure 1 (Section 1). For all decompositions, we observe that our algorithms obtain highly similar results within a few iterations for all the graphs. In particular, the average number of iterations we need in order to reach the 90% similarity are 5.4, 7.7 and 6 for k-core, k-truss, and $(3,4)$ nucleus decompositions, respectively. For the 99% similarity, 19.3, 17.7, and 12.5 iterations are needed on average. Most of the iterations after a certain point are performed just for a few r-cliques. We also observe that the first iteration, which uses the $\mathcal{S}$-degrees of r-cliques, provides

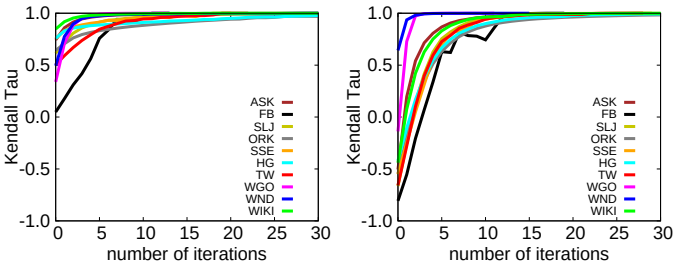

(a) k-core (b) k-truss

Figure 6: **Convergence rates on all the graphs. We measure the Kendall-Tau similarity between τ_i indices and κ_s indices at each iteration i. We obtain very close solutions to the exact decompositions within a few iterations. 90% similar results can be obtained with ∼80% less iterations than the exact decompositions, on average.**

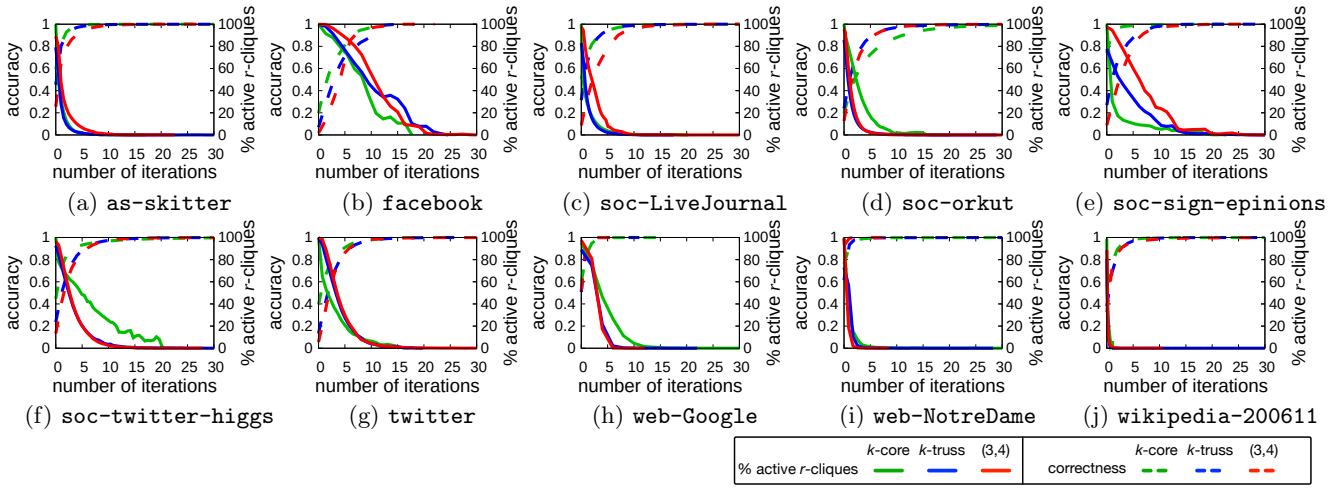

Figure 7: **Changes in the ratio of active r-cliques and the accuracy of τ indices during the computation, When the ratio of active r-cliques goes below 40%, τ indices provide 83%, 82% and 86% accurate results on average for k-core, k-truss and $(3,4)$ nucleus decompositions, respectively. If the ratio is below 10%, ∼99% accuracy is achieved in all decompositions.**

a great increase in the Kendall-Tau similarity score. This is especially significant for higher-order decompositions, k-truss and $(3,4)$. Overall, we can effectively approximate the κ_s indices in a few iterations.

In our algorithms, there is no way to infer if an r-clique has converged or not by checking the repetitions on τ indices at each iteration, because there might exist intermediate plateaus in τ indices, as explained in Section 4.2.1. However, we know which r-cliques are active or idle in each iteration. We check the ratio of active r-cliques during the computation, which also corresponds to the workload in each iteration, and examine their relations with the accuracy ratios of the r-cliques (what percentage has converged so far).

Figure 7 presents the results for each graph on all three decompositions. We observe that when the ratio of active r-cliques goes below 40% during the computation, 83%, 82%, and 86% accurate results are obtained for k-core, k-truss, and $(3,4)$ nucleus decompositions, on average. When the ratio goes below 10%, ∼99% accuracy is achieved in all decompositions, which is pretty significant. We believe that the ratio of active r-cliques is a helpful guide to find the sweet spots where almost-exact results are obtained more efficiently. Watching for 10 or 40 percentage of active r-cliques yields nice trade-offs between runtime and quality. If the computation is terminated when the ratio of active r-cliques goes below 10% and 40%, we observe up to 1.67x and 2.65x speedups with respect to full computation until convergence, in k-truss decomposition. For the $(3,4)$ nucleus decomposition, speedups can be as large as 1.29x and 2.27x if we stop the computation when the active r-cliques are below 10% and 40%. Overall, our algorithms enable trading-off between runtime and quality, and can be tailored to the needs of the real-world applications.

5.3 AND is significantly faster than the peeling

We compare the runtime performances of AND (Algorithm 3) and the peeling process (Algorithm 1) on three decompositions. We perform full computations until convergence. For the k-truss and $(3,4)$ nucleus decompositions, triangle counts per edge and four-clique counts per triangle need to be computed and we parallelize these parts for the peeling algorithms as well, for a fair comparison. Rest of the peeling

(in seconds)	k-core			k-truss			$(3,4)$		
Graphs	*Peeling*	AND alg.	*Speedup*	*Peeling*	AND alg.	*Speedup*	*Peeling*	AND alg.	*Speedup*
ASK	0.61	1.04	0.59	55.9	13.8	4.04	457	167	2.73
FB	0.002	0.008	0.21	0.533	0.6	0.84	36	21	1.70
SLJ	4.35	7.45	0.58	192.7	104.6	1.84	25111	11651	2.16
ORK	11.46	19.90	0.58	746.7	359.1	2.08	6781	2977	2.28
SSE	0.02	0.10	0.25	3.1	2.5	1.24	103	59	1.76
HG	0.70	1.10	0.64	158.1	73.0	2.17	2047	875	2.34
TW	0.04	0.07	0.50	5.3	4.6	1.15	131	73	1.79
WGO	0.36	0.36	0.99	7.2	2.8	2.61	45	14	3.28
WND	0.02	0.11	0.21	3.0	0.9	3.41	261	33	7.98
WIKI	3.84	9.36	0.41	378.3	96.5	3.92	3034	820	3.70

Table 5: **Runtime comparison of AND (Algorithm 3) and peeling (Algorithm 1) algorithms using 24 threads. Runtimes are in seconds. AND obtains significant speedups over the peeling process. In particular, $(3,4)$ is 7.98x faster than peeling on `web-NotreDame` graph, and k-truss decomposition has 4.04x speedup over the peeling process.**

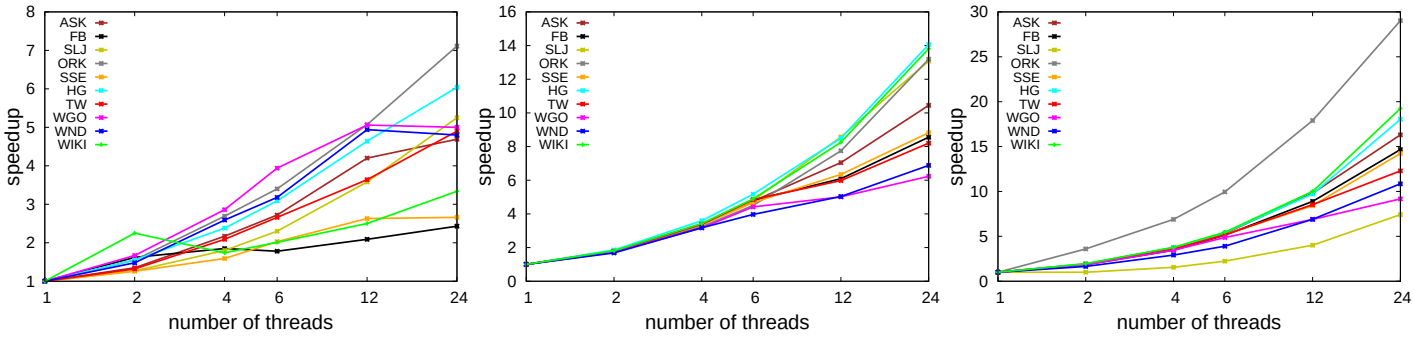

(a) k-core (b) k-truss (c) $(3,4)$

Figure 8: **Speedups of the parallel computations with respect to the sequential runs for the k-core, k-truss, and $(3,4)$ nucleus decompositions for all the graphs. We used $2,4,6,12$, and 24 threads where each thread is assigned to a single core. With 24 threads, we observe up to 7x, 14x, and 29x speedups for k-core, k-truss, and $(3,4)$ nucleus decompositions, respectively. Speedup numbers increase with more number of threads and faster solutions are possible with better machines that have more cores.**

computation is sequential as it cannot be parallelized. Runtimes with 24 threads are presented in Table 5. Overall, k-core computations are quite fast and there is not much work to benefit from parallelism, so the peeling process is faster than AND. Note that peeling performs a single scan of the graph while our algorithm does that at each iteration. For the k-truss and $(3,4)$ nucleus decompositions, AND obtains significant speedups over the peeling process. In particular, $(3,4)$ is 7.98x faster than peeling on the `web-NotreDame` graph, and k-truss decomposition has 4.04x speedup over the peeling process on the `as-skitter` graph.

We also check the speedups for the approximate decompositions. We show how the speedups (with respect to peeling algorithm) change when a certain amount of accuracy in κ_s indices is sacrificed. Figure 9 presents the behavior for k-truss and $(3,4)$ nucleus decompositions on some representative graphs, starting from iteration 0. We observe that after a certain accuracy score, speedups do not change significantly. For k-truss decomposition on `as-skitter` graph, accuracy scores jump from 0.45 to 0.8 in the first iteration and the speedup drops significantly – from 31 to 10. However, the accuracy scores above 0.8 do not change the speedup drastically. Slower rate of changes in speedup numbers (close to the convergence) enables to get more accurate results with a small additional cost. The reason behind this behavior is our notification mechanism that avoids the redundant computations. Similar trends are also observed for `soc-twitter-higgs` and `wikipedia-200611` graphs, on both k-truss and $(3,4)$ nucleus decompositions.

5.4 Shared-memory parallelization is scalable

We present the strong scalability evaluation for our shared-memory implementations of k-core, k-truss, and $(3,4)$ nucleus decompositions. Our machine has 24 cores in total (12 in each socket) and we run the experiments with $2,4,6,12$, and 24 threads where each thread is assigned to a single core. Speedups for all three decompositions with respect to the sequential runs of AND are given in Figure 8 for all the graphs. On average, using 24 threads we achieve 4.62x, 10.33x, and 15.12x speedups for k-core, k-truss, and $(3,4)$ nucleus decompositions, respectively. We also observe up to 7x, 14x, and 29x faster computations on some graphs. Our speedup numbers increase with more threads and *faster solutions are possible with more cores.* We also observe that speedups increase with larger r, s values which is due to the increasing computation demand. Overall, our parallel local algorithms are scalable and promise faster computations with more threads.

6. RELATED WORK

Previous attempts to estimate the core numbers (or k-cores) focus on the neighborhood of a vertex within a certain ra-

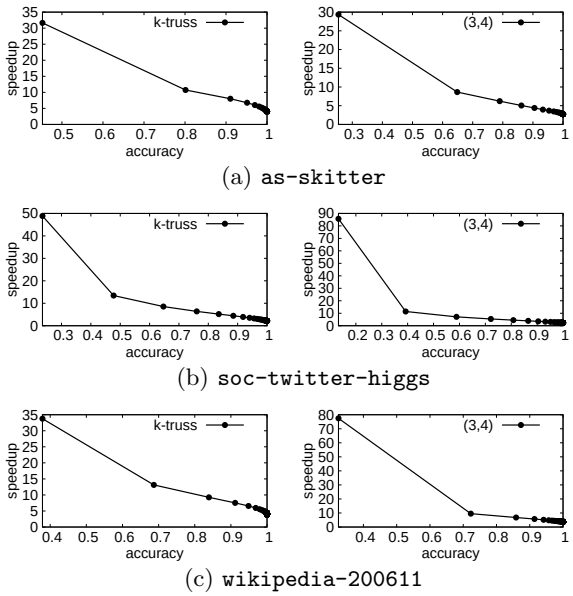

(a) `as-skitter`

(b) `soc-twitter-higgs`

(c) `wikipedia-200611`

Figure 9: **Change of speedups (over peeling algorithm) with varying approximations of k-truss, and $(3,4)$ nucleus decompositions. After a certain accuracy score, speedups do not change significantly, which enables to get more accurate results with a small additional cost. Thanks to our notification mechanism (Section 4.2.1), most of the work is done in the first few iterations and fine tuning of the κ_s indices are made easier.**

dius [22]. It is reported that if the radius is at least half of the diameter, close approximations can be obtained. However, given the small-world nature of the real-world networks, the local graph within a distance of half the diameter is too large to compute. In our work, we approximate the k-core, k-truss, and (r, s) nucleus decompositions in a rigorous and efficient way that does not depend on the diameter.

Most related study is done by Lu et al. [20], where they show that iterative h-index computation on vertices result in the core numbers. Their experiments on smaller graphs also show that h-index computation provides nice trade-offs for time and quality of the solutions. In our work, we generalized the iterative h-index computation approach for *any* nucleus decomposition that subsumes the k-core and k-truss algorithms. Furthermore, we give provable upper bounds on the number of iterations for convergence. Apart from that work, Govindan et al. [13] use the iterative h-index computation to design space-efficient algorithms for estimating core numbers.

Regarding the parallel computations, Jiang et al. [16] introduced parallel algorithms to find the number of iterations needed to find the empty k-core in random hypergraphs. Their nice work relies on the assumption that the edge density is below a certain threshold and their focus is on the number of iterations only. Our local algorithms present an alternative formulation for the peeling process, works for any k value. For the k-truss decomposition, Quick et al. [24] introduced algorithms for vertex-centric distributed graph processing systems. Their main focus is on the communication cost. For the same setup, Shao et al. [31] proposed faster algorithms that can compute k-trusses in a distributed graph processing system. Both papers make use of the peeling-based algorithms for computation. Our focus is on the generalization of the h-index based computation for the nucleus decompositions. Our local algorithms are pleasingly parallel.

7. CONCLUSION

We introduced a generalization of the iterative h-index computations for *any* nucleus decomposition and prove convergence bounds. Our local algorithms are pleasingly parallel and can provide approximations to explore time and quality trade-offs. Experimental evaluation on real-world networks exhibits the efficiency, scalability, and effectiveness of our algorithms for three decompositions. We believe that our parallel local algorithms for core, truss, and nucleus decompositions will be beneficial for many real-world applications.

Acknowledgements: This work was funded by the Applied Mathematics Program at the U.S. Department of Energy and Sandia's Laboratory Directed Research & Development (LDRD) program. Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

8. REFERENCES

[1] N. K. Ahmed, J. Neville, R. A. Rossi, and N. G. Duffield. Efficient graphlet counting for large networks. In *2015 IEEE International Conference on Data Mining, ICDM 2015, Atlantic City, NJ, USA, November 14-17, 2015*, pages 1–10, 2015.

[2] A. Angel, N. Sarkas, N. Koudas, and D. Srivastava. Dense subgraph maintenance under streaming edge weight updates for real-time story identification. *PVLDB*, 5(6):574–585, Feb. 2012.

[3] V. Batagelj and M. Zaversnik. An o(m) algorithm for cores decomposition of networks. Technical Report cs/0310049, Arxiv, 2003.

[4] A. R. Benson, D. F. Gleich, and J. Leskovec. Higher-order organization of complex networks. *Science*, 353(6295):163–166, 2016.

[5] J. Cohen. Trusses: Cohesive subgraphs for social network analysis. National Security Agency Technical Report, 2008.

[6] L. Dagum and R. Menon. Openmp: an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55, 1998.

[7] Y. Dourisboure, F. Geraci, and M. Pellegrini. Extraction and classification of dense communities in the web. In *WWW*, pages 461–470, 2007.

[8] X. Du, R. Jin, L. Ding, V. E. Lee, and J. H. T. Jr. Migration motif: a spatial - temporal pattern mining approach for financial markets. In *KDD*, pages 1135–1144, 2009.

[9] E. Fratkin, B. T. Naughton, D. L. Brutlag, and S. Batzoglou. Motifcut: regulatory motifs finding with maximum density subgraphs. In *ISMB*, pages 156–157, 2006.

[10] D. Gibson, R. Kumar, and A. Tomkins. Discovering large dense subgraphs in massive graphs. In *VLDB*, pages 721–732, 2005.

[11] A. Gionis, F. Junqueira, V. Leroy, M. Serafini, and I. Weber. Piggybacking on social networks. *PVLDB*, 6(6):409–420, 2013.

[12] D. F. Gleich and C. Seshadhri. Vertex neighborhoods, low conductance cuts, and good seeds for local community methods. In *KDD*, pages 597–605, 2012.

[13] P. Govindan, S. Soundarajan, T. Eliassi-Rad, and C. Faloutsos. Nimblecore: A space-efficient external memory algorithm for estimating core numbers. In *2016 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*, pages 207–214, Aug 2016.

[14] J. E. Hirsch. An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences of the United States of America*, 102(46):16569–16572, 2005.

[15] M. Jha, C. Seshadhri, and A. Pinar. Path sampling: A fast and provable method for estimating 4-vertex subgraph counts. In *Proceedings of the 24th International Conference on World Wide Web*, WWW '15, pages 495–505, 2015.

[16] J. Jiang, M. Mitzenmacher, and J. Thaler. Parallel peeling algorithms. In *Proceedings of the 26th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '14, pages 319–330, New York, NY, USA, 2014. ACM.

[17] R. Jin, Y. Xiang, N. Ruan, and D. Fuhry. 3-hop: a high-compression indexing scheme for reachability

query. In *SIGMOD Conf.*, pages 813–826, 2009.

[18] R. Kumar, P. Raghavan, S. Rajagopalan, and A. Tomkins. Trawling the web for emerging cyber-communities. In *WWW*, pages 1481–1493, 1999.

[19] J. Leskovec and A. Krevl. SNAP Datasets, June 2014.

[20] L. Lü, T. Zhou, Q.-m. Zhang, and H. E. Stanley. The h-index of a network node and its relation to degree and coreness. *Nature Communications*, 7:10168, 2016.

[21] D. Matula and L. Beck. Smallest-last ordering and clustering and graph coloring algorithms. *JACM*, 30(3):417–427, 1983.

[22] M. P. O'Brien and B. D. Sullivan. Locally estimating core numbers. In *ICDM*, pages 460–469, 2014.

[23] A. Pinar, C. Seshadhri, and V. Vishal. ESCAPE: efficiently counting all 5-vertex subgraphs. *CoRR*, abs/1610.09411, 2016.

[24] L. Quick, P. Wilkinson, and D. Hardcastle. Using pregel-like large scale graph processing frameworks for social network analysis. In *Proceedings of the 2012 International Conference on Advances in Social Networks Analysis and Mining (ASONAM 2012)*, ASONAM '12, pages 457–463, 2012.

[25] R. Rossi and N. Ahmed. The network data repository with interactive graph analysis and visualization. In *AAAI*, pages 4292–4293, 2015.

[26] R. A. Rossi, R. Zhou, and N. K. Ahmed. Estimation of graphlet statistics. *CoRR*, abs/1701.01772, 2017.

[27] K. Saito and T. Yamada. Extracting communities from complex networks by the k-dense method. In *ICDMW*, 2006.

[28] A. E. Sarıyüce, C. Seshadhri, A. Pınar, and Ü. V. Çatalyürek. Finding the hierarchy of dense subgraphs using nucleus decompositions. In *WWW*, pages 927–937, 2015.

[29] A. E. Sarıyüce, C. Seshadhri, A. Pınar, and Ü. V. Çatalyürek. Nucleus decompositions for identifying hierarchy of dense subgraphs. *ACM Transactions on Web (TWEB)*, to appear.

[30] S. B. Seidman. Network structure and minimum degree. *Social Networks*, 5(3):269–287, 1983.

[31] Y. Shao, L. Chen, and B. Cui. Efficient cohesive subgraphs detection in parallel. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*, SIGMOD '14, pages 613–624, New York, NY, USA, 2014. ACM.

[32] C. Tsourakakis. The k-clique densest subgraph problem. In *Proc. of the 24th International Conf. on World Wide Web*, WWW '15, pages 1122–1132, 2015.

[33] A. Verma and S. Butenko. Network clustering via clique relaxations: A community based approach. In *DIMACS Graph Part. and Clustering*, pages 129–140, 2012.

[34] Y. Zhang and S. Parthasarathy. Extracting analyzing and visualizing triangle k-core motifs within networks. In *ICDE*, pages 1049–1060, 2012.